

A goal-oriented RBM-Accelerated generalized polynomial chaos algorithm*

Jiahua Jiang[†] Yanlai Chen[†] Akil Narayan[‡]

September 19, 2016

Abstract

The non-intrusive generalized Polynomial Chaos (gPC) method is a popular computational approach for solving partial differential equations (PDEs) with random inputs. The main hurdle preventing its efficient direct application for high-dimensional input parameters is that the size of many parametric sampling meshes grows exponentially in the number of inputs (the “curse of dimensionality”). In this paper, we design a weighted version of the reduced basis method (RBM) for use in the non-intrusive gPC framework. We construct an RBM surrogate that can rigorously achieve a user-prescribed error tolerance, and ultimately is used to more efficiently compute a gPC approximation non-intrusively. The algorithm is capable of speeding up traditional non-intrusive gPC methods by orders of magnitude without degrading accuracy, assuming that the solution manifold has low Kolmogorov width. Numerical experiments on our test problems show that the relative efficiency improves as the parametric dimension increases, demonstrating the potential of the method in delaying the curse of dimensionality. Theoretical results as well as numerical evidence justify these findings.

1 Introduction

Computational methods for stochastic problems in uncertainty quantification (UQ) are an increasingly-important area of research and much recent effort in this direction has been rewarded with many promising developments. In particular, algorithms that quantify the effect of (potentially) random input parameters on solutions to differential equations have seen rapid advancement. One of the most widely used methods in this context is the generalized Polynomial Chaos (gPC) method [37], which constructs a parametric response surface (with parameters modeling the randomness) using a polynomial representation. This method exploits parametric regularity of the system to achieve fast convergence rates [35]. With gPC, stochastic solutions are represented as expansions in orthogonal polynomials of the input random parameters, and so many algorithms for parametric approximation concentrate on computation of the expansion coefficients in a gPC representation. Collocation-based or pseudospectral-based methods are popular non-intrusive approaches to compute these coefficients, using a collection of interpolation or quadrature nodes in parameter space [36]. This requires one to query an expensive yet deterministic computational solver once for each parameter node. However, when the dimension of the random parameter is large, the size of many sampling meshes (and hence the number of computational solves) grows exponentially. This is a manifestation of the “curse of dimensionality”.

*Submitted to the editors on 9/12/2016.

Funding: This work was funded by the National Science Foundation (NSF) under awards DMS-1216928 and DMS-1552238, the Air Force Office of Scientific Research (AFOSR) under award FA9550-15-1-0467, and the Defense Advanced Research Projects Agency (DAPRA) under award N660011524053.

[†]Mathematics Department, University of Massachusetts Dartmouth, North Dartmouth, MA (jjiang@umassd.edu, yanlai.chen@umassd.edu).

[‡]Department of Mathematics, and Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT (akil@sci.utah.edu).

One popular strategy that combats the computational burden arising from multiple queries of an expensive model is model order reduction, which includes proper orthogonal decomposition (POD) methods, Krylov subspace methods, and reduced basis methods (RBM). The references [6, 33] detail some of these methods. Model reduction strategies allow one to replace an expensive computational model with an inexpensive yet accurate emulator for which performing a large number of queries may be more computationally feasible.

Such an approach appeals to the same motivation as POD methods: although the discretized solution space is very high in dimension, the output of interest (such as the full solution field or integrated quantities of interest) frequently lie in a low-dimensional manifold. More precisely, the manifold can be approximated very well within a subspace of much lower dimension [7, 10, 27]. The search for, identification, and exploitation of this low-dimensional manifold are the central goals of many model order reduction strategies. Assuming such a low-dimensional manifold exists, then it may be possible to build a reduced-complexity emulator and consequently form the sought accurate gPC approximation in an efficient manner. In this paper we employ the RBM model reduction strategy for which [7, 18, 31] are good references with [2, 3, 5, 19] the appropriate historical references.

The Reduced Basis Method performs a projection onto a subspace spanned by “snapshots”, i.e., a small and carefully chosen selection of the most representative high-fidelity solutions. The fundamental reason this is an accurate approach is that, for many PDE’s of interest, the solution manifold induced by the parametric variation has small Kolmogorov width; see [25, 32] for a general discussion. These snapshots are selected via a weak greedy algorithm that appeals to an *a posteriori* error estimate [18, 31]. The computational methods that one uses to compute high-fidelity snapshots include typical solvers, like spectral collocation or finite element discretizations. The ingredient in RBM that allows for computational savings is the “offline-online” decomposition. The offline stage is the more expensive part of the algorithm where a small number (denoted N throughout this paper) of parameter values are chosen and the snapshots are generated by executing the expensive high-fidelity computational model at these parameter locations. The preparation completed during the offline stage allows very efficient evaluation of an emulator of the high-fidelity model at any other parameter value, i.e., “online”. During the online stage, each evaluation of the emulator can typically be computed orders of magnitude faster than evaluation of the original expensive model. The actual speedup depends largely on how fast the Kolmogorov width decays and how optimally the weak greedy algorithm [7] is able to select parameter values. Theoretical results on the quality of the N -dimensional surrogate space in approximating the full solution manifold are established in [10] and later improved in [7]. Roughly speaking, polynomial and exponential decay rates of the theoretical Kolmogorov width of the solution manifold (with respect to N) can be achieved by the algorithmic RBM procedure. In practice, one achieves 2 to 3 orders of magnitude speedup [13, 18]. One of the major benefits of RBM that we exploit in this paper is that the RBM model reduction is *rigorous*: Certifiable error bounds accompany construction of the emulator in the offline stage [31].

The idea of utilizing RBM for problems in a general uncertainty quantification framework is not new [8, 9, 11, 16, 21, 30]. In this paper our ultimate aim is to form a gPC expansion. The use of the RBM in this context, the design and analysis of an algorithm targeting statistical quantities of interest, and the exploration of its effectiveness in high-dimensional random space are underdeveloped to the best of our knowledge. A naïve stochastic collocation or pseudospectral method is computationally challenging in high-dimensional parameter spaces, even when employing a sparse grid of economical cardinality. The hybrid gPC-RBM algorithm we propose is able to reduce the computational complexity required for construction of a gPC surrogate, and assists construction of the gPC approximation in high-dimensional parameter spaces with rigorous error bounds on the gPC coefficients, and on any Lipschitz continuous quantity of interest of the resulting expansion.

This paper thus refines and extends the idea of combining a goal-oriented Reduced Basis Method with a generalized Polynomial Chaos expansion. Our framework is *goal-oriented*: the construction of the approximation is optimized with a user-specifiable quantity of interest in mind. The algorithm is *rigorous*: we can guarantee an error tolerance for general quantities of interest. Our numerical results indicate that the method improves in performance (efficiency) as the parametric dimension increases for the examples studied in this paper. This suggests that our method is particularly useful for delaying the curse of dimensionality.

The remainder of the paper is structured as follows. In section 2, we introduce the general framework of a PDE with parametric input data. The two major ingredients in our approach, gPC and RBM, are likewise discussed. In section 3 we introduce the novel contribution of this paper, the hybrid algorithm, which is analyzed in section 4. Our numerical results are collected in section 5, and verify the efficiency and convergence of the hybrid algorithm.

2 Background

In this section, we introduce the necessary background material of the hybrid algorithm, namely, generalized Polynomial Chaos and the Reduced Basis Method.

2.1 Problem setting

Let $\boldsymbol{\mu} = (\mu_1, \dots, \mu_K)$ be a K -variate random vector with mutually independent components on a complete probability space. For $\Gamma_i = \mu_i(\Omega)$ the image of μ_i , we denote the probability density function of the random variable μ_i as $\rho_i : \Gamma_i \rightarrow \mathbb{R}^+$. Since the components of $\boldsymbol{\mu}$ are mutually independent, then $\boldsymbol{\rho}(\boldsymbol{\mu}) = \prod_{i=1}^K \rho_i(\mu_i)$ is the joint probability density function of random vector $\boldsymbol{\mu}$. The image of $\boldsymbol{\mu}$ is $\boldsymbol{\Gamma} = \bigoplus_{i=1}^K \Gamma_i \subset \mathbb{R}^K$.

Let $D \subset \mathbb{R}^d$ ($d = 1, 2, 3$) be an open set in the physical domain with boundary ∂D , and $\mathbf{x} = (x_1, \dots, x_d) \in D$ a point in this set. We consider the problem of finding the solution $u : D \times \boldsymbol{\Gamma} \rightarrow \mathbb{R}$ of the following PDE with random parameters:

$$\begin{cases} \mathcal{L}(\mathbf{x}, u, \boldsymbol{\mu}) = f(\mathbf{x}, \boldsymbol{\mu}), & \forall (\mathbf{x}, \boldsymbol{\mu}) \in D \times \boldsymbol{\Gamma}, \\ \mathcal{B}(\mathbf{x}, u, \boldsymbol{\mu}) = g(\mathbf{x}, \boldsymbol{\mu}), & \forall (\mathbf{x}, \boldsymbol{\mu}) \in \partial D \times \boldsymbol{\Gamma}. \end{cases} \quad (2.1)$$

Here \mathcal{L} is a differential operator defined on domain D and \mathcal{B} is a boundary operator defined on the boundary ∂D . The functions f and g represent the forcing term and the boundary conditions, respectively.

We require the problem (2.1) to be well-posed and have a solution in a Hilbert space X . We thus assume $u(\cdot; \boldsymbol{\mu}) \in X$ for all $\boldsymbol{\mu}$. The Hilbert space X is equipped with inner product $(\cdot, \cdot)_X$ and induced norm $\|\cdot\|_X$. A canonical example is when (2.1) corresponds to a linear elliptic partial differential equation [17]: X satisfies $H_0^1(D) \subset X \subset H^1(D)$, with H^1 the space of functions whose first derivatives are square-integrable over D , and H_0^1 the space of functions in H^1 whose support is compact in D .

In most applications, one has access to a deterministic computational solver that, for each *fixed* value of $\boldsymbol{\mu}$, produces an approximate, discrete solution to (2.1). We assume that for this fixed $\boldsymbol{\mu}$, such a computational solver produces the discrete solution $u^\mathcal{N}$, which has \mathcal{N} degrees of freedom. This discrete solution is obtained by solving a discretized version of (2.1). For a fixed $\boldsymbol{\mu} \in \boldsymbol{\Gamma}$, this is given by

$$\begin{cases} \mathcal{L}^\mathcal{N}(u^\mathcal{N}, \boldsymbol{\mu}) = f^\mathcal{N}(\boldsymbol{\mu}), \\ \mathcal{B}^\mathcal{N}(u^\mathcal{N}, \boldsymbol{\mu}) = g^\mathcal{N}(\boldsymbol{\mu}). \end{cases} \quad (2.2)$$

Standard discretizations, such as finite element or spectral collocation solvers, can be written in this way. The continuous Hilbert space X is replaced with its discrete Hilbert space counterpart $X_\mathcal{N}$, with norm $\|\cdot\|_{X_\mathcal{N}}$.

As before, we assume that $u^\mathcal{N}(\boldsymbol{\mu}) \in X_\mathcal{N}$. We will need an additional assumption that the norm of the solution is uniformly bounded as a function of the parameter. I.e., that

$$\|u^\mathcal{N}(\cdot, \boldsymbol{\mu})\|_{X_\mathcal{N}} \leq U, \quad \forall \boldsymbol{\mu} \in \boldsymbol{\Gamma} \quad (2.3)$$

This assumption is satisfied for many practical problems of interest. For example, for a linear elliptic operator $\mathcal{L}^\mathcal{N}(u^\mathcal{N})$, boundedness of the solution is a simple consequence of the bilinear weak form

being coercive and the linear form being continuous [24]. In our setting, uniform coercivity of the bilinear form and uniform continuity of the linear form with respect to μ would be sufficient to guarantee the uniform boundedness (2.3).

When introducing the discretized PDE (2.2) we assume that, for any $\mu \in \Gamma$, $u^\mathcal{N}(\mu) \approx u(\mathbf{x}, \mu)$, where the approximation has an acceptable level of accuracy as determined by the modeling scenario. In practical modeling situations, one frequently requires a large number of degrees of freedom, $\mathcal{N} \gg 1$, to achieve this.

In what follows we will usually treat μ as a parameter associated with a ρ -weighted norm, rather than as an explicitly random quantity. This is a standard approach, and is without loss since all of our statements can be framed in the language of probability by appropriate change of notation. For example, the space of functions of μ with finite second moment is equivalent to the space $L_\rho^2(\Gamma)$,

$$\begin{aligned} L_\rho^2(\Gamma) &= \left\{ u : \Gamma \rightarrow \mathbb{R} \mid \int_\Gamma u^2(y) \rho(y) dy < \infty \right\} \\ &= \left\{ u : \Gamma \rightarrow \mathbb{R} \mid \mathbb{E} u^2(\mu) < \infty \right\}. \end{aligned}$$

2.2 Generalized Polynomial Chaos

The Generalized Polynomial Chaos method is a popular technique for solving stochastic PDE and representing stochastic processes [35]. The main idea of the gPC method is to seek an approximation of the exact solution of the PDE (2.1) by assuming that the dependence on μ is accurately represented by a finite-degree μ -polynomial. If u depends smoothly on μ , then exponential convergence with respect to the polynomial degree can be achieved. Computational implementations of gPC use an expansion in an orthogonal polynomial basis; as a consequence, quantities of interest such as expected value and variance can be efficiently evaluated directly from expansion coefficients.

2.2.1 gPC basis

Consider one-dimensional parameter space Γ_i corresponding to the random variable μ_i . If μ_i has finite moments of all orders, then there exists a collection of orthonormal polynomials $\left\{ \phi_m^{(i)}(\cdot) \right\}_{m=0}^\infty$, with ϕ_m a polynomial of degree m , such that

$$\mathbb{E} \left[\phi_m^{(i)}(\mu_i) \phi_n^{(i)}(\mu_i) \right] = \int_{\Gamma_i} \rho_i(\mu_i) \phi_m^{(i)}(\mu_i) \phi_n^{(i)}(\mu_i) d\mu_i = \delta_{m,n}$$

where $\delta_{m,n}$ is the Kronecker delta function. The type of orthogonal polynomial basis $\{\phi_m^{(i)}\}$ depends on the distribution of μ_i . For instance, if μ_i is uniformly distributed in $[-1, 1]$, its probability density function ρ_i is a constant and $\{\phi_m^{(i)}\}_{m=0}^\infty$ is the set of orthonormal Legendre polynomials. Several well-studied orthogonal polynomial families correspond to standard probability distributions [37].

For the K -dimensional case ($K > 1$), an orthonormal polynomial family associated to the full joint density $\rho(\mu)$ can be formed from products of univariate orthonormal polynomials:

$$\Phi_\alpha(\mu) = \phi_{\alpha_1}^{(1)}(\mu_1) \dots \phi_{\alpha_K}^{(K)}(\mu_K),$$

where $\alpha = (\alpha_1, \dots, \alpha_K) \in \mathbb{N}_0^K$ is a multi-index. The degree of Φ_α is $|\alpha| = \sum_{k=1}^K \alpha_k$. A standard polynomial space to consider in the multivariate setting is the *total degree* space, formed from the span of all Φ_α whose degree is less than a given $P \in \mathbb{N}$:

$$\mathcal{U}_K^P \equiv \text{span} \{ \Phi_\alpha \mid |\alpha| \leq P \} \quad (2.4)$$

The dimension of \mathcal{U}_K^P , denoted by M , is

$$M \equiv \dim(\mathcal{U}_K^P) = \binom{K+P}{K}, \quad (2.5)$$

which grows comparably to P^K for large K . In what follows, we will index multivariate orthonormal polynomials as either Φ_α with $\alpha \in \mathbb{N}_0^K$ satisfying $|\alpha| \leq P$, or Φ_m with $m \in \mathbb{N}$ satisfying $1 \leq m \leq \dim \mathcal{U}_K^P$. To achieve this, we assume any ordering of multi-indices α that preserves a partial ordering of the total degree (for example, graded lexicographic ordering).

2.2.2 gPC approximation and quadrature

The $L_\rho^2(\Gamma)$ -optimal gPC approximation of the solution $u(\mathbf{x}, \boldsymbol{\mu})$ to (2.1) in the space \mathcal{U}_K^P is the $L_\rho^2(\Gamma)$ -orthogonal projection onto \mathcal{U}_K^P , given by

$$u_K^P(\mathbf{x}, \boldsymbol{\mu}) = \sum_{m=1}^M \tilde{u}_m(\mathbf{x}) \Phi_m(\boldsymbol{\mu}). \quad (2.6)$$

The Fourier coefficient functions \tilde{u}_m are defined as

$$\tilde{u}_m(\mathbf{x}) = \int u(\mathbf{x}, \boldsymbol{\mu}) \Phi_m(\boldsymbol{\mu}) \rho(\boldsymbol{\mu}) d\boldsymbol{\mu}. \quad (2.7)$$

For any $\mathbf{x} \in D$, the mean-square error in this finite-order projection is

$$\begin{aligned} \mathcal{E}_{\text{gPC}}(\mathbf{x}) &= \|u(\mathbf{x}, \boldsymbol{\mu}) - u_K^P(\mathbf{x}, \boldsymbol{\mu})\|_{L_\rho^2(\Gamma)} \\ &= \left(\int_{\Gamma} (u(\mathbf{x}, \boldsymbol{\mu}) - u_K^P(\mathbf{x}, \boldsymbol{\mu}))^2 \rho(\boldsymbol{\mu}) d\boldsymbol{\mu} \right)^{1/2}. \end{aligned} \quad (2.8)$$

Note that this error is usually not achievable in practice: The Fourier coefficients \tilde{u}_m cannot be computed without essentially full knowledge of the solution u . Therefore, one frequently resorts to approximating these coefficients. One popular non-intrusive method is quadrature-based pseudospectral approximation, where the integral in (2.7) is approximated by a quadrature rule.

Toward that end, let $\{\boldsymbol{\mu}^q, w_q\}_{q=1}^Q$ denote quadrature nodes and weights, respectively, for a quadrature rule that implicitly defines a new empirical probability measure:

$$\int_{\Gamma} f(\boldsymbol{\mu}) \rho(\boldsymbol{\mu}) d\boldsymbol{\mu} \approx \sum_{q=1}^Q w_q f(\boldsymbol{\mu}^q). \quad (2.9)$$

For example, two common choices for quadrature rules are tensor-product Gauss quadrature rules, and Gauss-Patterson-based sparse grid quadrature rules (e.g., [15]). Each of these rules can integrate polynomials of high degree, but the requisite size of the quadrature rule Q is large in high dimensions. Both quadrature grids have cardinality that grows exponentially with dimension (for a fixed degree of integration accuracy), although sparse grids have a smaller size among the two. In this paper, we use tensor-product Gauss quadrature rules for low dimensions and sparse grid constructions for high dimensions; with these choices, we hereafter assume that (2.9) holds for functions f of the form of the integrand in (2.7).

With this quadrature rule, the Fourier coefficients can be approximated by

$$\tilde{u}_m \approx \hat{u}_m = \sum_{q=1}^Q u(\mathbf{x}, \boldsymbol{\mu}^q) \Phi_m(\boldsymbol{\mu}^q) w_q. \quad (2.10)$$

The advantage of this formulation is that we need only compute the quantities $u(\cdot, \boldsymbol{\mu}^q)$, which are a collection of solutions to a deterministic PDE. Since this is all done in the context of a computational solver given by (2.2), one will replace the continuous solution $u(\cdot, \boldsymbol{\mu})$ with the discrete solution $u^\mathcal{N}(\boldsymbol{\mu})$.

Then a straightforward stochastic quadrature approach first collects the solution ensemble from the computational solver, and subsequently uses it to compute approximate Fourier coefficients:

$$\{u_q^\mathcal{N}(\mathbf{x})\}_{q=1}^Q := \{u^\mathcal{N}(\mathbf{x}, \boldsymbol{\mu}^q)\}_{q=1}^Q \longrightarrow \hat{u}_m^\mathcal{N} = \sum_{q=1}^Q u_q^\mathcal{N}(\mathbf{x}) \Phi_m(\boldsymbol{\mu}^q) w_q, \quad (2.11)$$

The full approximation is formed by replacing the exact Fourier coefficients \tilde{u}_m in (2.6) with the $\hat{u}_m^{\mathcal{N}}$ coefficients computed above.

Note that, in order for the quadrature approximation (2.11) to be reasonably accurate, the number of quadrature points Q should be comparable with M . We already know from (2.5) that M scales like P^K for total-degree spaces. Typically, the cost of obtaining each $u_q^{\mathcal{N}}$ requires at least $\mathcal{O}(\mathcal{N})$ computational effort. (In some cases $\mathcal{O}(\mathcal{N}^3)$ effort is required.) Q solves of the PDE are required, with each solve costing at least $\mathcal{O}(\mathcal{N})$ work. Since $Q \sim M \sim P^K$, then in the best-case scenario the total work scales like $\mathcal{O}(\mathcal{N}P^K)$. Thus, the requisite computational effort for a straightforward stochastic quadrature method is onerous when K , the dimension of the random parameter $\boldsymbol{\mu}$, is large.

However, if these coefficients could be computed, then the resulting expansion (2.6) can be very accurate. The focus of this paper is to *inexpensively* compute an approximation to $\hat{u}_m^{\mathcal{N}}$. The essential ingredient is replacement of $u_q^{\mathcal{N}}$ by a surrogate that is much cheaper to compute, and whose approximation fidelity can be rigorously quantified.

2.2.3 Quantities of Interest

In many UQ scenarios, one is not necessarily interested in the entire solution field $u(\mathbf{x}, \boldsymbol{\mu})$, but rather some other quantity of interest derived from it. We introduce a functional \mathcal{F} that serves to map the solution u to the quantity of interest (the “goal”). Our construction exploits the well-known property of gPC that common quantities of interest such as the mean field and variance field can be recovered by simple manipulation of the gPC coefficients [35], up to the accuracy of the gPC expansion. Our theoretical results require two assumptions on the quantity of interest map \mathcal{F} .

The first assumption we make is that \mathcal{F} has affine dependence on an M -term gPC expansion, u_M , specifically

$$\mathcal{F}[u_M] = \mathcal{F}\left[\sum_{m=1}^M \hat{u}_m(\mathbf{x})\phi_m(\boldsymbol{\mu})\right] = \sum_{m=1}^M \theta_{\mathcal{F}}(\hat{u}_m(\mathbf{x}))\mathcal{F}[\phi_m(\boldsymbol{\mu})], \quad (2.12)$$

for some function $\theta_{\mathcal{F}}$. It is not hard to show that typical quantities of interest satisfy this condition on \mathcal{F} with simple coefficient functions $\theta_{\mathcal{F}}$:

- \mathcal{F} is the expected value operator \mathbb{E} and $\theta_{\mathcal{F}}$ is the identity function,

$$\mathcal{F}[u_M] = \mathbb{E}[u_M(\mathbf{x}, \boldsymbol{\mu})] = \sum_{m=1}^M \hat{u}_m(\mathbf{x})\mathbb{E}[\phi_m(\boldsymbol{\mu})] = \hat{u}_1(\mathbf{x}) \quad (2.13a)$$

- \mathcal{F} is the variance or norm-squared operator and $\theta_{\mathcal{F}}$ is the quadratic function $\theta_{\mathcal{F}}(v) = v^2$,

$$\mathcal{F}[u_M] = \text{var}(u_M(\mathbf{x}, \boldsymbol{\mu})) = \sum_{m=1}^M (\hat{u}_m(\mathbf{x}))^2 \text{var}[\phi_m(\boldsymbol{\mu})] = \sum_{m=2}^M (\hat{u}_m(\mathbf{x}))^2 \quad (2.13b)$$

$$\mathcal{F}[u_M] = \|u_M(\mathbf{x}, \boldsymbol{\mu})\|_{L_\rho^2}^2 = \sum_{m=1}^M (\hat{u}_m(\mathbf{x}))^2 \|\phi_m(\boldsymbol{\mu})\|_{L_\rho^2}^2 = \sum_{m=1}^M (\hat{u}_m(\mathbf{x}))^2. \quad (2.13c)$$

Our theoretical results also require a second assumption: that the functional $\theta_{\mathcal{F}}$ is Lipschitz continuous with Lipschitz constant C_{Lip} , i.e., that

$$\|\theta_{\mathcal{F}}(v) - \theta_{\mathcal{F}}(w)\| \leq C_{\text{Lip}} \|v - w\|, \quad (2.14)$$

for all appropriate inputs v and w . For $\mathcal{F} = \mathbb{E}$, this constant is $C_{\text{Lip}} = 1$. For the latter cases of $\mathcal{F} = \text{var}$ and $\mathcal{F}[\cdot] = \|\cdot\|_{L_\rho^2}^2$ where $\theta_{\mathcal{F}}(v) = v^2$, then $C_{\text{Lip}} = 2U$ with U is the uniform bound in (2.3). This is a consequence of $\|\theta_{\mathcal{F}}(v) - \theta_{\mathcal{F}}(w)\| = \|v^2 - w^2\| \leq \|v + w\| \cdot \|v - w\|$ and the uniform-boundedness of the solution (2.3).

2.3 Reduced Basis Method (RBM)

The reduced basis method [7, 18, 31] is a popular model order reduction strategy to solve a parameterized PDE for a large number of different parameter configurations. RBM seeks to form an approximation u^N satisfying

$$u^N(\boldsymbol{\mu}) \approx u^{\mathcal{N}}(\boldsymbol{\mu}), \quad \boldsymbol{\mu} \in \boldsymbol{\Gamma},$$

such that the approximation u^N can be computed with an algorithm whose complexity depends only on N , in contrast to the full solution $u^{\mathcal{N}}$ whose complexity depends on $\mathcal{N} \gg N$. To achieve this speedup, RBM algorithms traditionally make 3 assumptions:

- The Kolmogorov N -width of the (discretized) solution manifold is small, i.e.,

$$d_N := \inf_{\dim Y=N} \sup_{\boldsymbol{\mu} \in \boldsymbol{\Gamma}} \inf_{v \in Y} \|u^{\mathcal{N}}(\cdot, \boldsymbol{\mu}) - v\|_{X^{\mathcal{N}}} \quad (2.15)$$

is small when $N \ll \mathcal{N}$, where the outer infimum is taken over all N -dimensional subspaces of $X^{\mathcal{N}}$. In practice one hopes that d_N decays algebraically or even exponentially with increasing N . The small N -width requirement is a fundamental mathematical assumption without which RBM cannot achieve acceptable error with small N . RBM forms a dimension- N space X^N that seeks to achieve a distance to the solution manifold that is close to d_N .

- There is an efficiently-computable rigorous *a posteriori* error estimate. Given the N -dimensional RBM subspace X^N of $X^{\mathcal{N}}$, an estimate $\Delta_N(\boldsymbol{\mu})$ can be computed with $\mathcal{O}(N)$ complexity, satisfying

$$\Delta_N(\boldsymbol{\mu}) \geq \|u^{\mathcal{N}}(\cdot, \boldsymbol{\mu}) - u^N(\boldsymbol{\mu})\|_{X^{\mathcal{N}}}, \quad \boldsymbol{\mu} \in \boldsymbol{\Gamma} \quad (2.16)$$

The computable error estimate is required for an iterated, greedy construction of RBM approximation spaces X^N .

- The operators \mathcal{L} and f in (2.1) have affine dependence on $\boldsymbol{\mu}$. I.e., there exist $\boldsymbol{\mu}$ -independent operators \mathbb{L}_q and f_q , and \mathbf{x} -independent functions $\theta_q^L(\boldsymbol{\mu})$ and $\theta_q^f(\boldsymbol{\mu})$ such that

$$\mathcal{L}(\boldsymbol{\mu}) = \sum_{q=1}^{Q_L} \theta_q^L(\boldsymbol{\mu}) \mathbb{L}_q, \quad f(\mathbf{x}, \boldsymbol{\mu}) = \sum_{q=1}^{Q_f} \theta_q^f(\boldsymbol{\mu}) f_q(\mathbf{x}) \quad (2.17)$$

The affine assumptions are required so that RBM can compute the reduced-order solution u^N with $N \ll \mathcal{N}$ complexity.

In this paper, we will also assume that \mathcal{L} is a linear operator: this simplifies presentation of some RBM mechanics and allows easy motivation of formulas connecting PDE residuals with solution errors. While these assumptions are traditional, there are constructive remedies for non-affine and certain non-linear operators [4, 20, 29]. The RBM algorithm is a central part of the novel hybrid approach that we present in Section 3.

2.3.1 Reduced basis approximation

We recall from the discussion in Section 2.1 that a computational solver in (2.2) uses $\mathcal{N} \gg 1$ degrees of freedom to produce $u_q^{\mathcal{N}}(\mathbf{x})$, which is deemed an acceptably accurate approximation to $u(\mathbf{x}, \boldsymbol{\mu}^q)$. In the RBM context, this approximation is called the *truth solution* or *truth approximation* and we will use this terminology when appropriate. The starting point for developing computational reduced basis methods is to replace the expensive truth solution with an inexpensive reduced-order solution. We briefly describe the standard method for accomplishing this below.

Assume that a training set of parameter samples $\Xi \in \boldsymbol{\Gamma}$ is given such that the $\boldsymbol{\mu}$ -variation of the solution u is accurately captured by the resulting truth solution ensemble $\{u^{\mathcal{N}}(\mathbf{x}, \boldsymbol{\mu}) : \boldsymbol{\mu} \in \Xi\}$.

¹ This is a rather stringent requirement on Ξ , and does not furnish a constructive definition in general. In this paper we take Ξ to be the quadrature rule nodal set introduced in (2.9), that is, $\Xi = \{\boldsymbol{\mu}^q\}_{q=1}^Q$. Since our ultimate goal is formation of a gPC surrogate, we argue that such a choice of Ξ is a reasonable choice for the RBM training set.

For any given reduced-order dimension $N \ll \mathcal{N}$, we build the N -dimensional reduced basis space X^N by a greedy algorithm. This space is constructed as a span of “snapshots” (i.e., truth solutions) [26]

$$X^N = \text{span}\{u^N(\mathbf{x}, \boldsymbol{\nu}^1), \dots, u^N(\mathbf{x}, \boldsymbol{\nu}^N)\}, \quad (2.18)$$

where $\{\boldsymbol{\nu}^1, \dots, \boldsymbol{\nu}^N\} \subset \Xi$ are chosen in a greedy fashion. Assuming the greedy algorithm is performed in a sufficiently accurate fashion, then the dimension- N space approximates the manifold $u(\cdot, \boldsymbol{\Gamma})$ with an error comparable to the N -width [7]. The small N -width assumption (2.15) thus guarantees that a greedily-constructed X^N achieves a small approximation error. Existence of an efficiently-computable error estimate satisfying (2.16) allows a feasible greedy search for the $\boldsymbol{\nu}^k$ to be performed. (See sections 2.3.2 and 2.3.3 below.)

For a fixed $\boldsymbol{\mu}^* \in \boldsymbol{\Gamma}$, RBM approximates the truth solution $u^N(\mathbf{x}, \boldsymbol{\mu}^*)$ by an element from X^N . Thus, the RBM surrogate $u^N(\mathbf{x}, \boldsymbol{\mu}^*)$ can be represented as

$$u^N(\mathbf{x}, \boldsymbol{\mu}^*) = \sum_{k=1}^N c_k(\boldsymbol{\mu}^*) u^N(\mathbf{x}, \boldsymbol{\mu}^k) \quad (2.19)$$

RBM algorithms proceed by computing the $c_k(\boldsymbol{\mu}^*)$ so that the PDE residual with the ansatz (2.19) is as small as possible. The meaning of “small” is made precise by the prescription of an appropriate projection operator \mathcal{P} such that, using linearity of operator, the following holds:

$$\mathcal{P} \left[\sum_{k=1}^N c_k(\boldsymbol{\mu}^*) \mathcal{L}^N(\mathbf{x}, u^N(\mathbf{x}, \boldsymbol{\mu}^k), \boldsymbol{\mu}^*) \right] = \mathcal{P} [f^N(\mathbf{x}, \boldsymbol{\mu}^*)]. \quad (2.20)$$

Concrete examples of this abstract projection operator are the continuous L^2 projection onto X^N , a discrete ℓ^2 projection (least-squares) on the spatial mesh, or an empirical interpolation procedure [13, 14]. The affine assumption (2.17) allows computation of the $c_k(\boldsymbol{\mu}^*)$ with an N -dependent complexity (as opposed to $\mathcal{N} \gg N$ complexity): Using (2.17) in (2.20), we have

$$\mathcal{P} \left[\sum_{k=1}^N c_k(\boldsymbol{\mu}^*) \sum_{q=1}^{Q_L} \boldsymbol{\theta}_q^L(\boldsymbol{\mu}^*) \underline{\underline{\mathbb{L}_q(u^N(\mathbf{x}, \boldsymbol{\mu}^k))}} \right] = \mathcal{P} \left[\sum_{q=1}^{Q_f} \boldsymbol{\theta}_q^f(\boldsymbol{\mu}^*) \underline{\underline{f_q(\mathbf{x})}} \right] \quad (2.21)$$

We note that only the terms that are double-underlined require \mathcal{N} -dependent complexity to evaluate. However, these terms do *not* depend on $\boldsymbol{\mu}^*$, and so they may be computed and stored during the offline stage. We refer to [13, 34] for more details.

In the following sections we detail portions of the RBM algorithm that our novel algorithm amends: the error estimate and the greedy algorithm (sections 2.3.2 and 2.3.3, respectively).

2.3.2 *A posteriori* error estimate

The goal of this section is to compute a bound on the norm of $e(\mathbf{x}, \boldsymbol{\mu}) := u^N(\mathbf{x}, \boldsymbol{\mu}) - u^N(\mathbf{x}, \boldsymbol{\mu})$ without computing the truth solution u^N . Let $R_N : D \times \boldsymbol{\Gamma}$ denote the (Riesz representation of the) truth discretization residual from (2.2) using the reduced-order solution from X_N :

$$R_N(\mathbf{x}, \boldsymbol{\mu}) := f^N(\cdot; \boldsymbol{\mu}) - \mathcal{L}^N(u^N(\mathbf{x}, \boldsymbol{\mu}); \boldsymbol{\mu}) = \mathcal{L}^N(e(\mathbf{x}, \boldsymbol{\mu}), \boldsymbol{\mu}) \quad (2.22)$$

¹More precisely, we require that manifold of the finite ensemble of solutions $u^N(\cdot, \Xi)$ is an accurate surrogate for the full manifold $u^N(\cdot, \boldsymbol{\Gamma})$.

Let $\beta_{LB}(\boldsymbol{\mu})$ be a lower bound for the smallest eigenvalue of $\mathcal{L}^{\mathcal{N}}(\boldsymbol{\mu})^T \mathcal{L}^{\mathcal{N}}(\boldsymbol{\mu})$:

$$0 < \beta_{LB}(\boldsymbol{\mu}) \leq \min_v \frac{v^T (\mathcal{L}^{\mathcal{N}}(\boldsymbol{\mu}))^T \mathcal{L}^{\mathcal{N}}(\boldsymbol{\mu}) v}{\|v\|_{X_{\mathcal{N}}}} \quad (2.23)$$

Above, $\mathcal{L}^{\mathcal{N}}(\boldsymbol{\mu})$ should be understood as the matrix representation of the operator $\mathcal{L}^{\mathcal{N}}(\cdot; \boldsymbol{\mu})$. The relations (2.22) and (2.23) can be used to conclude [13]

$$\Delta_N(\boldsymbol{\mu}) := \frac{\|R_N(\mathbf{x}, \boldsymbol{\mu})\|_{X_{\mathcal{N}}}}{\sqrt{\beta_{LB}(\boldsymbol{\mu})}} \geq \|e(\mathbf{x}, \boldsymbol{\mu})\|_{X_{\mathcal{N}}} \quad \forall \boldsymbol{\mu} \in \Xi \quad (2.24)$$

Thus the *a posteriori* error estimate Δ_N is rigorous, satisfying (2.16). For the computation of this bound to be efficient, we must compute the residual R_N and β_{LB} in an \mathcal{N} -independent fashion. The residual can be computed with $\mathcal{O}(N)$ complexity by exploiting the same manipulations used in (2.21). The efficient evaluation of $\beta_{LB}(\boldsymbol{\mu})$ can be accomplished via the successive constraint linear optimization method (SCM) [12, 18, 22, 23] with the marginal computational cost for each $\boldsymbol{\mu}$ independent of the truth solution complexity \mathcal{N} . (There is a one-time \mathcal{N} -dependent overhead computation.)

With the ability to efficiently compute Δ_N , we can describe the greedy algorithm for choosing the RBM parameter snapshot locations $\{\boldsymbol{\nu}^k\}$.

2.3.3 Greedy Algorithm

Given a current set of parametric samples $\{\boldsymbol{\nu}^1, \dots, \boldsymbol{\nu}^k\}$ and the training parameter set Ξ , a new parameter value $\boldsymbol{\nu}^{k+1} \in \Xi$ is ideally selected as the $\boldsymbol{\mu} \in \Xi$ that maximizes the error between $u^{\mathcal{N}}(\cdot, \boldsymbol{\mu})$ and its projection onto X^k . Computationally, this error may be estimated by Δ_k :

$$\boldsymbol{\nu}^{k+1} = \operatorname{argmax}_{\boldsymbol{\mu} \in \Xi} \Delta_k(\boldsymbol{\mu}) \quad (2.25)$$

This process is repeated either until the maximum value of Δ_N is smaller than a user-prescribed tolerance ϵ_{tol} , or until N reaches a user-defined maximum value.

In total, only N queries of the truth solution are required. (One query for each parameter selected via (2.25).) However, each optimization (2.25) requires sweeping over Ξ , evaluating the error estimate Δ_k everywhere in the training set. Although we have described above that evaluation of Δ_k has a complexity that is only N -dependent, the cardinality of Ξ can be very large, and so this greedy optimization can still be expensive for large parametric dimensions.

We close this section by remarking that there are several algorithmic optimizations to speed up the computation in (2.25). For example, the evaluation of $\Delta_k(\boldsymbol{\mu})$ for each $\boldsymbol{\mu} \in \Xi$ is embarrassingly parallel. In addition, Δ_k is monotonically decreasing with k if the projection operator \mathcal{P} in (2.21) is an orthogonal projection. In this case, any $\boldsymbol{\mu} \in \Xi$ satisfying $\Delta_k(\boldsymbol{\mu}) < \epsilon_{\text{tol}}$ may be permanently removed from future greedy sweeps.

3 Hybrid Algorithm

Recall that straightforward stochastic quadrature requiring Q solves of (2.2), each having \mathcal{N} -dependent complexity, can be computationally burdensome when the random parameter dimension K is large. Our approach simply replaces $u^{\mathcal{N}}$ in (2.11) with an RBM surrogate u^N that is tailored toward gPC approximations. It thus ameliorates the cost-per-solve (reducing it to an N -dependent operation count) by using a gPC-goal-oriented variant of the RBM algorithm.

Our algorithm uses a modified *a posteriori* error estimate in the traditional RB greedy algorithm. This results in a RBM-gPC hybrid algorithm that can accurately construct a gPC surrogate more efficiently than standard pseudospectral methods.

Our convergence metric is the L^2_{ρ} norm, and so we modify the standard RBM algorithm described in section 2.3 so that rigorous L^2_{ρ} error estimates may be derived. We exploit the observation

that each $u(\mathbf{x}, \boldsymbol{\mu}^q)$ associated with parameter value $\boldsymbol{\mu}^q$ should have some quantitative measure of importance as indicated by the probability density $\rho(\boldsymbol{\mu}^q)$. This idea was explored in [30], but our version differs notably from earlier methods since we do not explicitly use ρ as a weight for the *a posteriori* error estimate.

3.1 Weighted *a posteriori* error estimate

Design of the RBM error estimate can significantly affect the performance of the resulting reduced basis method, particularly so for our goal-oriented approach. The approximate gPC coefficient formula (2.10) is the w_q -weighted inner product between the polynomial Φ_m and u^N . Thus, the error estimate should likewise be weighted using the quadrature weight w_q . We emphasize again that our strategy is different from using the probability density function ρ as done in [30]; even in simple one-dimensional cases, it is easy to see that $w_q \propto \rho(\boldsymbol{\mu}^q)$ (e.g., Gaussian quadrature or Clenshaw-Curtis quadrature).

We introduce the following weighted *a posteriori* error estimate $\Delta_N^w(\boldsymbol{\mu})$:

$$\Delta_N^w(\boldsymbol{\mu}^q) = \frac{\|R_N(\cdot, \boldsymbol{\mu}^q)\|_{X_N}}{\sqrt{\beta_{LB}(\boldsymbol{\mu}^q)}} \sqrt{Q|w_q|}, \quad q = 1, \dots, Q, \quad (3.1)$$

where $\beta_{LB}(\boldsymbol{\mu}^q)$ is a lower bound for the smallest eigenvalue of \mathcal{N}^N as given in (2.23), and R_N is the PDE residual of the order- N surrogate as defined in (2.22). Note that the novel quantity is the factor $\sqrt{Q|w_q|}$; since w_q corresponds to a Q -point normalized quadrature rule, the quantity $Q|w_q|$ frequently has $\mathcal{O}(1)$ magnitude. For example, a univariate Gauss quadrature rules on compact domains $Qw_q \sim \sqrt{1 - (\mu^q)^2} \rho(\mu^q) \sim \mathcal{O}(1)$ for fairly general ρ [28]. The absolute value bars in (3.1) are necessary in general because sparse grid quadrature rules can have negative weights. For a tensor-product Gaussian quadrature rule, the weights are all positive.

The m -th Fourier coefficient produced by the hybrid algorithm is \hat{u}_m^N in (2.11) and its surrogate \hat{u}_m^N

$$\hat{u}_m^N = \sum_{q=1}^Q u^N(\boldsymbol{\mu}^q) \Phi_m(\boldsymbol{\mu}^q) w_q. \quad (3.2)$$

The form of the weighted *a posteriori* error estimate ensures that we can bound the error between this coefficient and the corresponding truth Fourier coefficient. The precise estimate appears in section 4.

3.2 Goal-oriented greedy algorithm

The greedy algorithm strategy here is essentially the same as in Section 2.3.3. One major difference is that we replace the original error estimate with our weighted one:

$$\boldsymbol{\nu}^{k+1} = \underset{\boldsymbol{\mu} \in \Xi}{\operatorname{argmax}} \Delta_k^w(\boldsymbol{\mu}) \quad (3.3)$$

We show later in Theorem 3 and Corollary 4 that the weighting provided by Δ_k^w allows one to guarantee that the gPC approximation that is formed from the RBM surrogates is within a user-defined tolerance of the gPC-truth approximation.

Another major difference in the goal-oriented algorithm is that the tolerance criterion is tuned to the quantity of interest of the gPC surrogate. At each stage with k RB snapshots, we compute the error estimate

$$\varepsilon = C_{Q,M} \sqrt{\frac{1}{Q} \sum_{q=1}^Q (\Delta_k^w)^2(\boldsymbol{\mu}^q)}, \quad \text{with } C_{Q,M} = \sum_{m=1}^M \sqrt{\sum_{q=1}^Q |w_q| \Phi_m^2(\boldsymbol{\mu}^q) |\mathcal{F}[\Phi_m(\boldsymbol{\mu})]|}, \quad (3.4)$$

where \mathcal{F} denotes the quantity of interest as introduced in (2.12). Note that the constant $C_{Q,M}$ is computable *independent* of the solution u , and depends only on the choice of quadrature rule and quantity of interest. (See Lemma 7, and the discussion following Corollary 4.) As we show in Corollary 4, ε is an upper bound on the error in the quantity of interest defined by \mathcal{F} between the inexpensive RBM surrogate and the expensive gPC truth approximation.

A pseudocode implementation of the weighted greedy approach is shown in Algorithm 1.

Algorithm 1 Goal-oriented greedy algorithm

- 1: Input: training set Ξ with associated quadrature weights w_q ;
 - 2: Input: stopping criterion tolerance ε_{tol} ;
 - 3: Input: goal-oriented constant $C_{Q,M} = \sum_{m=1}^M B_{Q,m} \mathcal{F}[\Phi_m(\boldsymbol{\mu})]$.
 - 4: Randomly select the first sample $\boldsymbol{\mu}^1 \in \Xi$;
 - 5: Obtain truth solution $u^{\mathcal{N}}(\mathbf{x}, \boldsymbol{\mu}^1)$, set $X_1 = \text{span}\{u^{\mathcal{N}}(\mathbf{x}, \boldsymbol{\mu}^1)\}$;
 - 6: Set $k = 1$ and $\varepsilon = \infty$;
 - 7: **while** $\varepsilon > \varepsilon_{\text{tol}}$ **do**
 - 8: **for** each $\boldsymbol{\mu} \in \Xi$ **do**
 - 9: Obtain RBM solution $u^k(\mathbf{x}, \boldsymbol{\mu})$ by computing $c_j(\mathbf{x})$ that satisfy (2.20)
 - 10: Compute weighted *a posteriori* error estimate $\Delta_k^w(\boldsymbol{\mu})$ from (3.1)
 - 11: **end for**
 - 12: Choose $\boldsymbol{\mu}^{k+1} = \text{argmax}_{(\boldsymbol{\mu} \in \Xi)} \Delta_k^w(\boldsymbol{\mu})$;
 - 13: augment the reduced basis space $X_{k+1} = \text{span}\{X_k \cup \{u^{\mathcal{N}}(\mathbf{x}, \boldsymbol{\mu}^{k+1})\}\}$;
 - 14: Calculate $\Delta^{\text{sum}} = \sum_{\boldsymbol{\mu} \in \Xi} (\Delta_N^w)^2(\boldsymbol{\mu})$;
 - 15: Set $\varepsilon = C_{Q,M} \sqrt{\frac{1}{|\Xi|} \Delta^{\text{sum}}}$;
 - 16: Set $N \leftarrow N + 1$;
 - 17: **end while**
-

3.3 Goal-oriented hybridized RBM-gPC algorithm

Summary pseudocode of the goal-oriented gPC-RBM procedure is shown in Algorithm 2.

The offline stage is Algorithm 1: the RBM approximation space is built by scanning the quadrature node set Ξ and evaluating the weighted *a posteriori* error estimate Δ_k^w . The expensive PDE solver (2.2) is queried a total of N times over the greedily constructed parameter set $\boldsymbol{\nu}^1, \dots, \boldsymbol{\nu}^N$. The offline phase stops when the computed error indicator ε defined in (3.4) falls below the user-defined tolerance ε_{tol} .

The online stage proceeds as indicated on line 5 of Algorithm 2. In this phase, the RBM surrogate that was constructed in the offline phase is evaluated several times to compute the approximate gPC coefficients $\{\hat{u}_m^N\}_{m=1}^M$. This portion of the algorithm is *much* faster than naïve evaluation of (2.11); it is in this section of the procedure where the hybrid algorithm reaps computational savings compared to a traditional stochastic pseudospectral approach.

Once the approximate coefficients \hat{u}_m^N are collected, the quantity of interest $\mathcal{F}[u_M^N]$ may be evaluated. A significant benefit of using the hybrid approach is that the error in this computed quantity of interest can be rigorously controlled by the user-defined input tolerance ε_{tol} ; we show this in the next section.

Thus, the hybrid algorithm *both* achieves significant computational savings in construction of a gPC approximation *and* provides strict error bounds on quantities of interest.

Algorithm 2 Hybrid Algorithm

- 1: Input: the general stochastic PDE (2.1) and a tolerance ε_{tol} .
- 2: **Offline procedure:**
- 3: Set Ξ as the training set and use the goal-oriented greedy procedure with tolerance ε_{tol} (Algorithm 1) to compute N reduced basis elements $\{u^N(\mathbf{x}, \boldsymbol{\mu}^j)\}_{j=1}^N$.
- 4: **end Offline procedure**
- 5: Using the online RBM surrogate $u^N(\mathbf{x}, \boldsymbol{\mu})$, compute M -term gPC coefficients and approximation of u^N :

$$\hat{u}_m^N = \sum_{q=1}^Q w_q u^N(\mathbf{x}, \boldsymbol{\mu}^q) \Phi_m(\boldsymbol{\mu}^q)$$

$$u_M^N(\mathbf{x}, \boldsymbol{\mu}) = \sum_{m=1}^M \hat{u}_m^N(\mathbf{x}) \Phi_m(\boldsymbol{\mu}).$$

- 6: Output: Solution u_M^N , or quantity of interest $\mathcal{F}[u_M^N] = \sum_{m=1}^M \theta_{\mathcal{F}}[\hat{u}_m^N(\mathbf{x})] \mathcal{F}[\Phi_m(\boldsymbol{\mu})]$.
-

3.3.1 Computational cost

The computational complexities of each step of the hybrid algorithm are as follows. Let $\mathcal{S}(\mathcal{N})$ denote the cost of one truth solve; depending on the problem and the solver employed, this typically varies between $\mathcal{O}(\mathcal{N})$ and $\mathcal{O}(\mathcal{N}^3)$.

The offline portion of the algorithm requires repeated scanning of parameter space for maximization of Δ_k^w . A naïve scanning requires Q evaluations of Δ_k^w . However, we can be more efficient. For example, we may trim entries in the set $\{\boldsymbol{\mu} \in \Xi : \Delta_N^w(\boldsymbol{\mu}) < \frac{\varepsilon_{\text{tol}}}{2C_{Q,M}}\}$ from Ξ after each loop. We let $n(\Xi, k)$ denote the size of the set

$$n(\Xi, k) = \left| \left\{ \boldsymbol{\mu} \in \Xi \mid \Delta_{k-1}^w \leq \frac{\varepsilon_{\text{tol}}}{2C_{Q,M}} \right\} \right|, \quad \Delta_0^w(\boldsymbol{\mu}) \equiv \infty$$

Obviously $n(\Xi, k) \leq |\Xi| \equiv Q$, and $n(\Xi, k+1) \leq n(\Xi, k)$. Finally, we define $Q_N^{\text{trim}} := \sum_{k=1}^N n(\Xi, k)$. Using this notation we can give a rough operation count of the entire hybrid gPC-RBM algorithm.

The offline portion of the algorithm has complexity of the order

$$\underbrace{QM + Q_f^2 \mathcal{N}}_{\text{offline preparation}} + \underbrace{(N^4 + Q_L^2 N^3 + Q_f Q_L N^2 + Q_f^2) Q_N^{\text{trim}}}_{\text{Evaluation of } \Delta_k^w \text{ over } \Xi} + \underbrace{N \mathcal{S}(\mathcal{N})}_{\text{truth solver}} + \underbrace{(Q_L^2 N^2 + Q_f Q_L) \mathcal{N}^2}_{\text{online preparation}}$$

Below are more detailed explanations of this operation count:

- *Offline preparation* — evaluation of $C_{Q,M}$ and computations enabling optimized evaluations of $\|f(\boldsymbol{\mu})\|_{\mathcal{X}^N}$.
- *Evaluation of Δ_k^w* — $n(\Xi, k)$ evaluations for $k = 1, \dots, N$ results in the Q_N^{trim} factor. Each evaluation at stage k requires $k^2 Q_L + k Q_f + k^3$ operations to evaluate the k th RBM surrogate and its PDE residual, plus $Q_f^2 + Q_L^2 k^2 + Q_f Q_L k$ operations to evaluate the error estimate.
- *truth solver* — N solves of the truth-discretized PDE (2.2)
- *online preparation* — for each $k = 1, \dots, N$, optimizations for fast evaluations of Δ_{k+1}^w require $Q_L^2 k \mathcal{N}^2 + Q_f Q_L \mathcal{N}^2$, and optimizations for efficient computation of the k th RB surrogate require $Q_L k \mathcal{N}^2 + Q_f \mathcal{N}$ operations. The sum of these is dominated by the term shown.

The online portion of the algorithm has complexity on the order of $(N^2 Q_L + N Q_f + N^3) Q$. We emphasize that this complexity is \mathcal{N} -independent, but it does depend on Q .

Remark 1. *From this analysis, we see that it is worthwhile to apply the hybrid algorithm only when $QS(\mathcal{N})$ (the brute-force cost of (2.11)) is more than the total offline and online time combined. This is typically achievable when the N -width is small (so that $N \ll \mathcal{N}$) and when $Q_f Q_L \ll \mathcal{N}$. On the other hand, we note that there are extreme scenarios (such as when Q_L , Q_f , and N are all large) when it is more efficient to compute gPC coefficients directly via (2.11).*

The hybrid algorithm cost still scales linearly with Q , with $Q \sim P^K$ for tensorial quadrature rules. For this reason, the algorithm only delays the curse of dimensionality. However, when the RBM procedure is effective, this algorithm is still orders-of-magnitude more efficient than a direct gPC approximation via stochastic pseudospectral approximation.

4 Analysis of the hybrid algorithm

In this section, we show the convergence of this goal-oriented gPC-RBM algorithm. More precisely, we show that the error committed by the hybrid RBM algorithm is controlled by the user-defined input parameter ε_{tol} in Algorithm 2. Given a P -th order M -term gPC projection (2.6), the m -th truth Fourier coefficient $\hat{u}_m^{\mathcal{N}}$ is evaluated by (2.11) and its surrogate \hat{u}_m^N by (3.2). The truth and reduced basis stochastic solutions are then, respectively,

$$u_M^{\mathcal{N}}(\mathbf{x}, \boldsymbol{\mu}) = \sum_{m=1}^M \hat{u}_m^{\mathcal{N}}(\mathbf{x}) \Phi_m(\boldsymbol{\mu}), \quad (4.1)$$

$$u_M^N(\mathbf{x}, \boldsymbol{\mu}) = \sum_{m=1}^M \hat{u}_m^N(\mathbf{x}) \Phi_m(\boldsymbol{\mu}). \quad (4.2)$$

The properties of the RBM algorithm allow us to bound the error between the computationally expensive $\mathcal{F}[u_M^{\mathcal{N}}]$ and the efficiently computable hybrid surrogate $\mathcal{F}[u_M^N]$. This bound is our main theoretical result and is shown in Theorem 3.

To begin, we first need to control the error between the full stochastic quadrature gPC coefficients $u_m^{\mathcal{N}}(\mathbf{x})$ and its RBM surrogate $u_m^N(\mathbf{x})$.

Lemma 2. *Given a Q -point quadrature rule $\{\boldsymbol{\mu}^q, w_q\}_{q=1}^Q$ and an N -dimensional reduced basis approximation $u^N(\mathbf{x}, \boldsymbol{\mu})$ for the solution $u(\mathbf{x}, \boldsymbol{\mu})$ to the PDE (2.1), the error in the m -th gPC coefficient is given by*

$$\|\hat{u}_m^{\mathcal{N}}(\mathbf{x}) - \hat{u}_m^N(\mathbf{x})\|_{X_{\mathcal{N}}} \leq B_{Q,m} \sqrt{\frac{1}{Q} \sum_{q=1}^Q (\Delta_N^w)^2(\boldsymbol{\mu}^q)},$$

where $B_{Q,m}$ is the uncentered second moment of $\Phi_m(\boldsymbol{\mu})$ under the discrete measure defined by the quadrature rule:

$$B_{Q,m} = \sqrt{\sum_{q=1}^Q |w_q| \Phi_m^2(\boldsymbol{\mu}^q)}. \quad (4.3)$$

Proof. We use the quadrature representation for these functions to prove the result:

$$\begin{aligned}
\|\hat{u}_m^{\mathcal{N}}(\mathbf{x}) - \hat{u}_m^N(\mathbf{x})\|_{X_{\mathcal{N}}} &= \left\| \sum_{q=1}^Q w_q \Phi_m(\boldsymbol{\mu}^q) [u^{\mathcal{N}}(\mathbf{x}, \boldsymbol{\mu}^q) - u^N(\mathbf{x}, \boldsymbol{\mu}^q)] \right\|_{X_{\mathcal{N}}} \\
&\leq \sum_{q=1}^Q \left| \sqrt{|w_q|} \Phi_m(\boldsymbol{\mu}^q) \right| \left(\sqrt{|w_q|} \|e(\mathbf{x}, \boldsymbol{\mu}^q)\|_{X_{\mathcal{N}}} \right) \\
&\leq \sqrt{\sum_{q=1}^Q |w_q| \Phi_m^2(\boldsymbol{\mu}^q)} \sqrt{\sum_{q=1}^Q |w_q| \|e(\mathbf{x}, \boldsymbol{\mu}^q)\|_{X_{\mathcal{N}}}^2} \\
&\leq B_{Q,m} \sqrt{\sum_{q=1}^Q \frac{1}{Q} (\Delta_N^w)^2(\boldsymbol{\mu}^q)}
\end{aligned}$$

□

This result bounds the error in each gPC coefficient as a product of two terms: the first term $B_{Q,m}$ is a measure of the (polynomial degree of) accuracy of the quadrature rule, which is computable and independent of the solution u . The second term is an average of the weighted *a posteriori* error estimate over parameter space.

We can now bound the RBM error in the quantity of interest.

Theorem 3. *Given an M -term gPC projection (2.6) and an N -dimensional reduced basis approximation (2.20), the error in the quantity of interest computed from the RBM-gPC approximation u_M^N , and that computed from the truth gPC approximation $u_M^{\mathcal{N}}$ is*

$$\|\mathcal{F}[u_M^{\mathcal{N}}] - \mathcal{F}[u_M^N]\|_{X_{\mathcal{N}}} \leq C_{\text{Lip}} C_{Q,M} \sqrt{\frac{1}{Q} \sum_{q=1}^Q [\Delta_N^w(\boldsymbol{\mu}^q)]^2}, \quad (4.4)$$

where C_{Lip} is the Lipschitz constant of \mathcal{F} defined in (2.14), and $C_{Q,M}$ is a constant independent of u , defined by

$$C_{Q,M} = \sum_{m=1}^M B_{Q,m} |\mathcal{F}[\Phi_m(\boldsymbol{\mu})]|, \quad (4.5)$$

with $B_{Q,m}$ defined in (4.3).

Proof. We begin by using our assumption (2.12) regarding the affine dependence of \mathcal{F} on a gPC representation:

$$\mathcal{F}[u_M^{\mathcal{N}}] - \mathcal{F}[u_M^N] = \sum_{m=1}^M (\theta_{\mathcal{F}}(\hat{u}_m^{\mathcal{N}}(\mathbf{x})) - \theta_{\mathcal{F}}(\hat{u}_m^N(\mathbf{x}))) \mathcal{F}[\Phi_m(\boldsymbol{\mu})].$$

Applying triangle inequality and Lipschitz continuity of $\theta_{\mathcal{F}}(\cdot)$, we have

$$\|\mathcal{F}[u_M^{\mathcal{N}}] - \mathcal{F}[u_M^N]\|_{X_{\mathcal{N}}} \leq C_{\text{Lip}} \sum_{m=1}^M \|\hat{u}_m^{\mathcal{N}}(\mathbf{x}) - \hat{u}_m^N(\mathbf{x})\|_{X_{\mathcal{N}}} |\mathcal{F}[\Phi_m(\boldsymbol{\mu})]|.$$

Using Lemma 2 on the right-hand side, we obtain

$$\|\mathcal{F}[u_M^{\mathcal{N}}] - \mathcal{F}[u_M^N]\|_{X_{\mathcal{N}}} \leq C_{\text{Lip}} \sum_{m=1}^M B_{Q,m} |\mathcal{F}[\Phi_m(\boldsymbol{\mu})]| \sqrt{\frac{1}{Q} \sum_{q=1}^Q (\Delta_N^w)^2(\boldsymbol{\mu}^q)},$$

and this proves (4.4). □

Corollary 4. *The output gPC approximation u_M^N from Algorithm 2 satisfies*

$$\|\mathcal{F}[u_M^N] - \mathcal{F}[u_M^N]\|_{X_N} \leq C_{\text{Lip}} \varepsilon$$

where ε is defined in (3.4). In particular, ε is dominated by the user-prescribed ϵ_{tol} if Algorithm 2 terminates successfully.

Remark 5. *The Lipschitz constant C_{Lip} is trivially 1 when $\mathcal{F}[\cdot]$ is the expected value operator. For the other two cases listed in (2.13b) and (2.13c), it is finite as long as we have uniform stability with respect to the parameter $\boldsymbol{\mu}$ for the computational solver (2.2). As the discussion around (2.3) indicates, this is a standard assumption, and in that case $C_{\text{Lip}} = 2U$.*

Remark 6. *We emphasize that $C_{Q,M}$ is a scalar whose value is independent of the solution u , its truth discretization u^N , or the RBM surrogate u^N . It depends only on the choice of quadrature rule, the gPC order, and the choice of what quantity of interest is to be computed.*

The coefficient $C_{Q,M}$, depending on $B_{Q,m}$, is not analytically computable in general since it depends on the chosen quadrature rule in parameter space. However, the next lemma shows that for one of the two quadrature rules we are using in this paper, the tensor-product Gauss quadrature rule, we have $B_{Q,m} \equiv 1$ if the accuracy is sufficiently high.

Lemma 7. *Let $\{\Phi_1, \dots, \Phi_M\}$ span the degree- P isotropic total-degree space, so that $M = \binom{K+P}{P}$. Assume that the quadrature rule (2.9) corresponds to an isotropic tensor product quadrature rule with q points in each dimension, totaling $Q = q^K$ nodes. If $P < q$, then $B_{Q,m} \equiv 1$ for $m = 1 \dots, M$.*

Proof. This is a simple consequence of the fact that a q -point Gaussian quadrature rule exactly integrates polynomials up to degree $2q - 1$. With α the size- K multi-index corresponding to the linear index m , then

$$\Phi_m^2(\boldsymbol{\mu}) = \prod_{k=1}^K \phi_{\alpha_k}^2(\mu_k), \quad \alpha_k \leq |\alpha| = P.$$

Thus, in dimension k , we have $\deg \phi_{\alpha_k}^2 \leq 2P < 2q$. Additionally, a Gauss quadrature rule has all positive weights. Therefore, the quadrature rule integrates the polynomial in each dimension exactly, so $B_{Q,m}^2 = \sum_{j=1}^Q w_j \Phi_m^2(\boldsymbol{\mu}^j) = \mathbb{E} \Phi_m^2(\boldsymbol{\mu}) = 1$. \square

Of course, similar statements about $B_{Q,m}$ can be made for non-total-degree or anisotropic spaces so long as one has a good understanding of the quadrature rule. Even if one cannot analytically derive values for $B_{Q,m}$, it is easily and inexpensively computable by applying the quadrature rule to the gPC basis Φ_m^2 .

Thus under certain assumptions, the constants $C_{Q,M}$ and $B_{Q,m}$ are explicitly computable. We summarize some of these results in Table 1. In the case of a sparse grid quadrature rule, we cannot compute these constants analytically, but they can be computed numerically with ease using available sparse grid software (e.g., [1]).

As an illustration, consider dimension $K = 4$ with two different tensorial gPC basis sets: Legendre polynomials, and Jacobi polynomials (with parameters $\alpha = 1, \beta = 1$). We use a Gauss-Patterson-based sparse grid and compute $B_{Q,m}$. The results are shown in Figure 1. It is difficult to discern a pattern for the constants shown in the figure, but they are all $\mathcal{O}(1)$ for the range of parameters shown.

5 Numerical results

In this section, we present numerical results to illustrate the accuracy of the proposed hybrid approach and its efficiency compared to the conventional gPC method. The PDE with random inputs

Table 1: Summary of explicit values for bounding constants $B_{Q,m}$ and $C_{Q,M}$ for common quantity of interest operators \mathcal{F} . In the table, $\delta_{j,k}$ is the Kronecker delta function.

Quantity of interest	$B_{Q,m}$	$C_{Q,M}$	Assumptions
Mean value, $\mathcal{F} = \mathbb{E}$	1	1	The quadrature rule (2.9) exactly integrates the constant function, all weights are positive
Variance, $\mathcal{F} = \text{var}$	$\sqrt{\sum_{q=1}^Q w_q \Phi_m^2(\mu_q)}$	$\sum_{m=2}^M \sqrt{\sum_{q=1}^Q w_q \Phi_m^2(\mu_q)}$	None
	$1 - \delta_{m,1}$	$M - 1$	Assumptions of Lemma 7
L^2_ρ -norm squared, $\mathcal{F}[\cdot] = \ \cdot\ _{L^2_\rho}^2$	$\sqrt{\sum_{q=1}^Q w_q \Phi_m^2(\mu_q)}$	$\sum_{m=1}^M \sqrt{\sum_{q=1}^Q w_q \Phi_m^2(\mu_q)}$	None
	1	M	Assumptions of Lemma 7




Figure 1: Value of $B_{Q,m}$ with Legendre polynomial (left) and Jacobi polynomial (right). Q varies with the accuracy level of sparse grid (level 7 to level 16). m records the sequence of gPC bases with $degree \leq 5$

is the following linear elliptic equation posed on the spatial domain $D = [-1, 1] \times [-1, 1]$ with homogeneous Dirichlet boundary conditions,

$$\begin{cases} -\nabla \cdot (a(\mathbf{x}, \boldsymbol{\mu}) \nabla u(\mathbf{x}, \boldsymbol{\mu})) = f & \text{in } D \times \Gamma, \\ u(\mathbf{x}, \boldsymbol{\mu}) = 0 & \text{on } \partial D \times \Gamma. \end{cases} \quad (5.1)$$

The diffusion coefficient $a(\mathbf{x}, \boldsymbol{\mu})$ is defined as:

$$a(\mathbf{x}, \boldsymbol{\mu}) = A + \sum_{k=1}^K \frac{\cos(30\mu_k - 1)}{k^2} \cos(kx) \sin(ky),$$

where K is the parameter dimension and we set A to be a positive constant that is large enough so that the equation is uniformly elliptic on D . We take as the right hand side $f = 1$. As a first step, we test our algorithm on this canonical example [17] which is often tested in the gPC community [36,37]. As we mentioned in (2.12), the output of interest \mathcal{F} is defined as a certain functional of the solution over the physical domain D . Here, we explore the following two cases:

- $\mathcal{F} = \mathbb{E}$. By (2.13a), the error of the mean value we observe is

$$\xi_{\text{mean}} = \|\mathcal{F}[u_M^N] - \mathcal{F}[u_M^N]\|_{\ell^2(D)} = \|\hat{u}_1^N(\mathbf{x}) - \hat{u}_1^N(\mathbf{x})\|_{\ell^2(D)} \quad (5.2)$$

- $\mathcal{F} = \|\cdot\|_{L_\rho^2}^2$. By (2.13c), the error of norm-squared operator we evaluate is

$$\xi_{\text{norm}} = \|\mathcal{F}[u_M^{\mathcal{N}}] - \mathcal{F}[u_M^N]\|_{\ell^2(D)} = \left\| \sum_{m=1}^M (\hat{u}_m^{\mathcal{N}}(\mathbf{x}))^2 - \sum_{m=1}^M (\hat{u}_m^N(\mathbf{x}))^2 \right\|_{\ell^2(D)} \quad (5.3)$$

In our numerical experiments, we test problem (5.1) with $A = 5$ and $K = 2, 4, 6$. For the gPC approximation, we use the degree-5 total degree space, \mathcal{U}_K^5 defined in (2.4). We adopt a tensor product quadrature rule for lower dimensional case $K = 2$, and a Gauss-Patterson-based sparse grid for the higher dimensional cases $K = 4, 6$. See Table 2 for the nodal count Q and dimension M of the gPC approximation space for each K . The truth approximation is the solution from a pseudospectral solver on a $\mathcal{N} = 35 \times 35$ spatial grid. The online solver of the goal-oriented reduced basis method (i.e., the projection operator \mathcal{P} defined in (2.20)) is the least squares reduced collocation method developed in [13]. We test the algorithm for two different probability distributions for the random variable $\boldsymbol{\mu}$, namely, the uniform distribution and Beta distribution with shape parameters $\alpha = 2, \beta = 2$.

We plot the expected value $\mathbb{E}(u_M^{\mathcal{N}})$ in Figure 2. The error estimates ε defined in (3.4) with $\mathcal{F} = \mathbb{E}$ and $\mathcal{F} = \|\cdot\|_{L_\rho^2}^2$ are plotted in Figure 3 and Figure 4 respectively, both displaying exponential convergence. We then evaluate the actual error in the quantity of interest of the resulting surrogate solution, as defined in (5.2) and (5.3). These are shown in Figure 5 for the two choices, $\mathbb{E}(\cdot)$ and $\|\cdot\|_{L_\rho^2}^2$. We note that both ξ_{mean} and ξ_{norm} have a clear exponential trend in convergence for both probability distributions. Finally, we measure the efficiency of the hybrid algorithm by calculating the ratio of the runtime between those of the hybrid and the traditional gPC quadrature approach. This is plotted in Figure 6. Here the time for the hybrid algorithm includes *both* offline and online time. It is clear that the proposed hybrid gPC-RBM method can reach a high level of accuracy (Figure 5) while significantly alleviating the computational burden (Figure 6). Moreover, we observe that the efficiency is increasing as K gets larger, at least for the type of equations that is currently tested. Therefore, this alleviation is more significant for high-dimensional problems, indicating great potential of the hybrid approach for larger parametric dimensions.

Table 2: The number of quadrature nodes Q and gPC approximation space dimension M as a function of parametric dimension K .

K	2	4	6
Q(K)	1,600	22,401	367,041
M(K)	21	126	462

6 Conclusion

We propose, analyze, and numerically test a hybridized RBM-gPC algorithm. It is based on a newly designed weighted RBM enabling a particular greedy algorithm tailored for any applicable quantity of interest in the context of uncertainty quantification. The final algorithm is analyzed to be reliable, and tested to be accurate. We observe that the efficiency of the hybrid algorithm increases with respect to the dimension of parameter in the partial differential equation. This suggests that the hybrid approach may be useful in alleviating the curse of dimensionality for other problem as well.

References

- [1] *TASMANIAN software suite*, <http://tasmanian.ornl.gov/>.




Figure 2: Expected value when $K = 4$ for uniform (left) and Beta (right) distribution.




Figure 3: Error estimate ε with $\mathcal{F} = \mathbb{E}$ for uniform distribution (left) and Beta distribution (right).

- [2] A.K.NOOR AND J.M.PETERS, *Reduced basis technique for nonlinear analysis of structures*, AIAA, 18 (1980), pp. 455–462.
- [3] B. O. ALMROTH, P. STERN, AND F. A. BROGAN, *Automatic choice of global shape functions in structural analysis*, AIAA Journal, 16 (1978), pp. 525–528.
- [4] M. BARRAULT, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *An empirical interpolation method: application to efficient reduced-basis discretization of partial differential equations*, Comptes Rendus Mathématique, 339 (2004), pp. 667–672, doi:10.1016/j.crma.2004.08.006, <http://www.sciencedirect.com/science/article/pii/S1631073X04004248> (accessed 2012-10-05).
- [5] A. BARRETT AND G. REDDIEN, *On the reduced basis method*, Z. Angew. Math. Mech., 75 (1995), pp. 543–549.
- [6] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems*, SIAM Review, 57 (2015), pp. 483–531, doi:10.1137/130932715, <http://epubs.siam.org/doi/abs/10.1137/130932715> (accessed 2015-11-18).
- [7] P. BINEV, A.COHEN, W. DAHMEN, R. DEVORE, G. PETROVA, AND P. WOJTASZCZYK, *Convergence rates for greedy algorithms in reduced basis methods*, SIAM Journal of Mathematical Analysis, 43 (2011), pp. 1457–1472, doi:10.1137/100795772.




Figure 4: Error estimate ε with $\mathcal{F} = \|\cdot\|_{L^2}^2$ for uniform distribution (left) and Beta distribution (right).

- [8] S. BOYAVAL, C. L. BRIS, T. LELIÈVRE, Y. MADAY, N. NGUYEN, AND A. PATERA, *Reduced basis techniques for stochastic problems*, Archives of Computational Methods in Engineering, 17 (2010), pp. 435–454, doi:10.1007/s11831-010-9056-z.
- [9] S. BOYAVAL, C. LE BRIS, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *A reduced basis approach for variational problems with stochastic parameters: application to heat conduction with variable Robin coefficient*, Comput. Methods Appl. Mech. Engrg., 198 (2009), pp. 3187–3206, doi:10.1016/j.cma.2009.05.019, <http://dx.doi.org/10.1016/j.cma.2009.05.019>.
- [10] A. BUFFA, Y. MADAY, A. T. PATERA, C. PRUDHOMME, AND G. TURINICI, *A priori convergence of the Greedy algorithm for the parametrized reduced basis method*, ESAIM: Mathematical Modelling and Numerical Analysis, 46 (2012), pp. 595–603, doi:10.1051/m2an/2011056.
- [11] P. CHEN AND A. QUARTERONI, *A new algorithm for high-dimensional uncertainty quantification based on dimension-adaptive sparse grid approximation and reduced basis methods*, Journal of Computational Physics, 298 (2013), pp. 176–193, doi:10.1016/j.jcp.2015.06.006.
- [12] Y. CHEN, *A certified natural-norm successive constraint method for parametric inf-sup lower bounds*, Applied Numer. Math., 99 (2016), pp. 98–108.
- [13] Y. CHEN AND S. GOTTLIEB, *Reduced Collocation Methods: Reduced Basis Methods in the Collocation Framework*, Journal of Scientific Computing, 55 (2012), pp. 718–737, doi:10.1007/s10915-012-9654-z.
- [14] Y. CHEN, S. GOTTLIEB, AND Y. MADAY, *Parametric analytical preconditioning and its applications to the reduced collocation methods.*, C. R. Acad. Sci. Paris, Ser. I, 352 (2014), pp. 661 – 666, doi:10.1016/j.crma.2014.06.001.
- [15] P. CONRAD AND Y. MARZOUK, *Adaptive smolyak pseudospectral approximations*, 35, pp. A2643–A2670, doi:10.1137/120890715.
- [16] H. C. ELMAN AND Q. LIAO, *Reduced basis collocation methods for partial differential equations with random coefficients*, SIAM/ASA J. Uncertainty Quantification, 1 (2013), pp. 192–217, doi:10.1137/120881841.
- [17] L. C. EVANS, *Partial differential equations*, vol. 19 of Graduate Studies in Mathematics, American Mathematical Society, Providence, RI, 1998.




Figure 5: Error of the expected value (top) and norm-squared operator (bottom) for uniform distribution (left) and Beta distribution (right).

- [18] D. H. G. ROZZA AND A. PATERA, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: Application to transport and continuum mechanics*, Archives of Computational Methods in Engineering, 15 (2008), pp. 229–275, doi:10.1007/s11831-008-9019-9.
- [19] M. GREPL AND A. PATERA, *A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations*, ESAIM: Mathematical Modelling and Numerical Analysis, 39 (2005), pp. 157–181, doi:http://dx.doi.org/10.1051/m2an:2005006.
- [20] M. A. GREPL, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations*, ESAIM: Mathematical Modelling and Numerical Analysis, 41 (2007), pp. 575–605, doi:10.1051/m2an:2007031.
- [21] B. HAASDONK, K. URBAN, AND B. WIELAND, *Reduced basis methods for parameterized partial differential equations with stochastic influences using the Karhunen-Loève expansion*, SIAM/ASA J. Uncertain. Quantif., 1 (2013), pp. 79–105, doi:10.1137/120876745, <http://dx.doi.org/10.1137/120876745>.
- [22] D. HUYNH, D. KNEZEVIC, Y. CHEN, J. HESTHAVEN, AND A. PATERA, *A natural-norm successive constraint method for inf-sup lower bounds*, CMAME, 199 (2010), pp. 1963–1975.
- [23] D. HUYNH, G. ROZZA, S. SEN, AND A. PATERA, *A successive constraint linear optimization method for lower bounds of parametric coercivity and inf-sup stability constants*, C. R. Acad. Sci. Paris, Série I., 345 (2007), pp. 473 – 478.




Figure 6: Efficiency of the hybrid algorithm for uniform distribution (left) and Beta distribution (right).

- [24] C. JOHNSON, *Numerical solution of partial differential equations by the finite element method*, Cambridge University Press, Cambridge, 1987.
- [25] G. G. LORENTZ, M. V. GOLITSCHKE, AND Y. MAKOVZ, *Constructive Approximation - Advanced Problems*, no. 304 in Grundlehren der mathematischen Wissenschaften, Springer, 1996.
- [26] Y. MADAY, *Reduced basis method for the rapid and reliable solution of partial differential equations*, in International Congress of Mathematicians. Vol. III, Eur. Math. Soc., Zürich, 2006, pp. 1255–1270.
- [27] Y. MADAY, A. T. PATERA, AND G. TURINICI, *A Priori Convergence Theory for Reduced-Basis Approximations of Single-Parameter Elliptic Partial Differential Equations*, Journal of Scientific Computing, 17 (2002), pp. 437–446, doi:10.1023/A:1015145924517, <http://link.springer.com/article/10.1023/A%3A1015145924517> (accessed 2013-07-16).
- [28] P. NEVAI, T. ERDLYI, AND A. MAGNUS, *Generalized Jacobi Weights, Christoffel Functions, and Jacobi Polynomials*, SIAM Journal on Mathematical Analysis, 25 (1994), pp. 602–614, doi:10.1137/S0036141092236863.
- [29] N. C. NGUYEN, A. T. PATERA, AND J. PERAIRE, *A best points interpolation method for efficient approximation of parametrized functions*, International Journal for Numerical Methods in Engineering, 73 (2008), pp. 521–543, doi:10.1002/nme.2086, <http://onlinelibrary.wiley.com/doi/10.1002/nme.2086/abstract> (accessed 2015-10-29).
- [30] A. Q. P. CHEN AND G. ROZZA, *A weighted reduced basis method for elliptic partial differential equations with random input data*, SIAM Journal on Numerical Analysis, 51 (2013), pp. 3163–3185, doi:10.1137/130905253.
- [31] A. PATERA AND G. ROZZA, *Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations*, Copyright MIT, (2007), pp. 229–275, <http://augustine.mit.edu>.
- [32] A. PINKUS, *N-widths in approximation theory*, Springer, 1985.
- [33] A. QUARTERONI AND G. ROZZA, eds., *Reduced Order Methods for Modeling and Computational Reduction*, Springer International Publishing, <http://link.springer.com/10.1007/978-3-319-02090-7>.

- [34] G. ROZZA, D. B. P. HUYNH, AND A. T. PATERA, *Reduced Basis Approximation and a Posteriori Error Estimation for Affinely Parametrized Elliptic Coercive Partial Differential Equations*, Archives of Computational Methods in Engineering, 15 (2008), pp. 229–275, doi:10.1007/s11831-008-9019-9, <http://link.springer.com/article/10.1007/s11831-008-9019-9> (accessed 2014-09-22).
- [35] D. XIU, *Fast Numerical Methods for Stochastic Computations: A Review*, Communications in Computational Physics, 5 (2009), pp. 242–272, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.148.5499> (accessed 2010-04-10).
- [36] D. XIU AND J. S. HESTHAVEN, *High-Order Collocation Methods for Differential Equations with Random Inputs*, SIAM Journal on Scientific Computing, 27 (2005), pp. 1118–1139, doi:10.1137/040615201, <http://link.aip.org/link/?SCE/27/1118/1> (accessed 2010-04-10).
- [37] D. XIU AND G. E. KARNIADAKIS, *The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations*, SIAM Journal on Scientific Computing, 24 (2002), pp. 619–644, doi:10.1137/S1064827501387826, <http://link.aip.org/link/?SCE/24/619/1> (accessed 2009-12-01).